platform solutions
on-line news for developers

Issue 17, February 26, 1999

Feature Story
Each month we provide a feature article on key industry trends and developments. Authored by a member of Intel's Executive Staff, it offers insightful and useful information for product development, planning and execution.

Inside Looking In
Senior Technical Marketing Manager for Platform Technologies Tim Mostad gives you a fresh perspective on the latest technologies making their way onto Intel Architecture platforms. Tim lets you see the work through the eyes of the people making it happen and lets you hear what they really think. It's straight talk from developers to developers.

Pentium® III Processor Platform Series
Pentium® III Processor Platform Series PSN will bring you focused articles on the hottest technologies about to arrive on Pentium III processor-based platforms. From Rambus memory technology to optimizing software, it's news and information you can definitely use, and you'll find it every month here in PSN.

Top News Stories
Delivering an in-depth report on key platforms, products and technologies, our Top Stories provide a monthly source of information on the issues affecting hardware developers. Be sure to check in every month for the latest stories that are driving the evolution of the industry.

Reader Services
If you are new to *Platform Solutions* and would like to receive this companion newsletter to the on-line version, please visit *Platform Solutions* on-line and go the "Subscribe Now" section to register and sign up for delivery. The on-line version provides lots of direct links for quick access to the developer information and news reported in each issue, whether it's on Intel's Web site or industry Web sites. Please visit the following URL: http://developer.intel.com/solutions/

We want you to consider *Platform Solutions* as your personal information resource for the Intel® Architecture platform. If you can help us make it better, or if you have a comment, question or a specific topic you would like to see covered, we want to hear from you. Please take the opportunity to send us an e-mail with your specific feedback or request to: platform.solutions@intel.com

If you do not want to receive this mailing in the future, please send an email to: platform.solutions@intel.com with "unsubscribe" in the body of the message.

On behalf of all of us at Platform Solutions, welcome to the future of the PC platform!

Table of Contents

(Click on descriptions to jump to articles)

Feature

The Internet: Growth Driver for the Industry

by Paul Otellini
Executive Vice President and General Manager
Intel® Architecture Business Group, Intel Corporation

The Internet is driving the growth of the industry—toward a billion connected PCs, millions of servers and annual e-commerce revenues that could top a trillion dollars a year by the year 2002. Intel Executive Vice President Paul Otellini describes how the Pentium® III processor and other advances in the Intel® Architecture are building the foundation for Internet-driven market growth across all computing segments.

The Internet has emerged as the single most important growth driver for the computing industry. The Internet is not only attracting new PC users, it is also expanding the usage model for people who currently own PCs. The magnitude of the Internet phenomenon is aptly quantified in Intel's vision of a billion Internet-connected PCs and millions of Internet servers, and the economic implications for the industry are staggering. As the Internet becomes more pervasive, it is projected that annual e-commerce industry revenues will top one trillion dollars.

Today, the growth of the Internet has made computing ubiquitous. E-mail is already a universal business tool. Voice and data applications are converging over Internet Protocol-based network infrastructures. From the largest corporations to individual consumers, electronic commerce is changing our business models and the way we live. In the process, enterprise computing is becoming synonymous with Internet-connected computing. The enterprise, in turn, will continue to provide the Intel® Architecture-based infrastructure to support the Internet.

The growth of the Internet, the ubiquity of computing, and continuing advances in Intel Architecture are also proving the success of Intel's fundamental PC market segmentation strategy. The Internet runs on Intel Architecture, and Intel products offer solutions for all segments of the market. New growth opportunities continue to emerge across all segments.

In the performance desktop market segment, the new Pentium® III processor is the first Intel processor specifically designed to empower a new Internet experience, with Internet Streaming SIMD Extensions enriching the user experience with crisper video, state-of-the-art 3D and more vivid audio.

The success of Intel's market segmentation strategy was further underscored in the value PC segment by last year's Celeron™ processor introduction. The Celeron processor experienced the fastest ramp of any new product in Intel's history—a five-fold growth from the second quarter to the end of 1998.

In the mobile PC segment, Intel's 0.18-micron process technology is driving the emergence of new products and new levels of performance. This transition is now enabling the mobile PC to become a primary computing platform many users—as mobile systems continue to move ever-closer to desktop-equivalent performance and overall benefits of ownership.

Intel is now the performance leader in the workstation and the volume server market segments, while maintaining its long-standing leadership in price/performance.

In every computing segment, Intel will continue to support the industry by supplying the technologies and the high-performance products needed to attract new users and drive the adoption of new Internet-based hardware and software. It is a truism in our business that technology can mask complexity. It does so by enabling the industry to deliver enhanced products with greater Ease of Use to ensure greater user satisfaction.

The Internet is transforming the way the industry does business, and the way we do business at Intel. As we look ahead to 1999, Intel will continue to provide technology leadership the industry can rely on to enable, empower and enrich the Internet experience.

About the Author

Paul Otellini is responsible for Intel Architecture businesses and strategies. Intel business units that report to him include the Business Platform Group, the Consumer Products Group, the Mobile and Handheld Products Group and the Enterprise Server Group.

Inside Looking In

The Zen of Board Repair and DVD Standards Setting

By Tim Mostad
Senior Technical Marketing Manager
Intel Corporation

This month, Tim Mostad gives us an inside look at the sometimes mysterious and often illogical world of "standards setting"—a process that may not always follow the best technical implementations, but should make life simpler for developers.

Some 18 years ago, I started my Intel career as a repair technician, straight out of electronics school. It was a small step from the process of book learning to a world where all the answers were right in front of you if you just knew where to look. One thing you could rely on was the certainty that a product had worked at one time and getting to an answer didn't require the least bit of politics or business tradeoffs. Nor did I need to have the least bit of appreciation of the roles these two arts had in my employment.

As I would sit there, playing microelectronic sleuth, freeze spray liberally applied to board and staring at one big mystery, my job was so complex and yet so simple. Unwinding a couple of dozen of sheets of schematics, trying to diagnose abhorrent behavior with a scope, logic analyzer, current probe, or whatever, was not easy, but at least there was an answer you could understand

When I "graduated" out of the repair center after a few years and entered the world of product development, my schooling, in fact, resumed. I didn't design the things that broke; I just made sure they were serviceable. Even so, I got a good look at the rather arcane process of product design. Being a studious technical person, a.k.a. a geek, I thought the designer's job was simple: always do what makes technical sense. I quickly learned the attributes of supplier quality, component delivery, PCB quality, plating technology, package pin-out, etc., etc. have an importance of their own. At times, it seemed the designer was so hamstrung that there was no room for what made technical sense.

However, evidently, products eventually did make it out the door. Why? Only because of the decisions that the designer *didn't* have to make.

For all the choices a designer has to make, an equal number don't have to be made because designers can buy standard components off the shelf and put them into standard form factors. These components and form factors are based on standards set forth by someone else, developed on machines based on yet other standards. It's a whole chain of agreements and de facto decisions dating to antiquity. It's a good thing, too, or nothing would ever get done.

So just how do engineers create standards in the first place? Well, I've discovered they don't—at least not anymore (if they ever did). After my post graduate degree in product design, I have so closely approached insanity that I've qualified for a job helping to enable the adoption of new technology standards. I was very fortunate to experience the establishment of the PCI Local Bus Interface Standard from the inside as my first foray into this kooky world. PCI was standard setting done right.

To foster the adoption of PCI, Intel bellied up and licensed a number of very valuable patents on the condition that anyone else participating agreed to do the same. They did, and the collaboration was marvelous. The crazy part is that this accord would not have been possible without some very astute attorneys and managers willing to take significant business risks. The truth is that anyone can invent things. Those things become valuable only when people use them.

The PCI process paid off with an incredibly successful product ramp with millions of users and an interface standard so successful that it is everywhere. We immediately somehow rationalized that the process was no good and have not done it the same way since. So it may be true that “success breeds success” but our current attempts are something like “second cousins” to what we have done in the past. So can we still be successful? Our platform technology track record is pretty good with USB, AGP, and many others but it gets a lot tougher as we try to grow beyond our core competencies. Our current effort on DVD is a good test of our abilities to see what we can accomplish as we push the envelope.

DVD standards are driven mainly by the consumer electronics companies, many of the same ones that brought us Betamax, VHS, S VHS, VHS-C, 8 mm, Hi-8, and other formats for video. So many bright ideas exist that consumers were going to have the pleasure of sorting out DVD standards as well. Left to fate, each of us would get to choose between DVD-RAM, DVD-RW, DVD+RW (or as it is sometimes called just +RW) and/or DVD-ROM.

This is not exactly what we have been accustomed to in the PC world. While PCI had a facsimile of a competitor, there was only one and it faded quickly. We PC developers felt we could bring our “standards setting acumen” to consumer electronics and help make sense out of the format chaos.

To Intel, the picture should be clear. All DVD drives should read all DVD formats, as well as CD-ROM. The DVD writers can write whatever format just as long as it’s one of the formats all drives can read. This approach is simple enough in theory but it assumes the greatest business risk. Each company has to be willing to share for the greater good, and more daunting yet, their attorneys have to be able to agree. So how does an insignificant outsider like Intel make our vision real? As we learned from PCI, collaborate.

We found another very interested and influential computer company to act as an ally to help develop a proposal. Our initial proposal and theirs didn’t match. In the process of negotiation, we abandoned our position and accepted theirs, and (sigh) they adopted our now-abandoned position and left the one we just accepted. However, this did not represent a stalemate where players swapped opposing positions. Instead, it was a prime example of what happens when technical people just want to agree on something that makes sense: it often happens a bit too easily. The attorneys and managers in other businesses would thoroughly argue their position for years and prevent these kinds of “mistakes” from happening.

Ultimately, we have managed to come back to a plan we can both agree upon, and our newer, bigger, better plan has been unveiled to the DVD Forum. Now the real work begins as we try to get the plan adopted.

And as much as we technical folks would like, this process will not be logical and it won’t follow anything we learned in school. What’s more, in the end the result may not be anywhere close to any plan we’ve helped develop. The point of standards, though, is not to be logical or even represent the best technical implementation, but to simplify the lives of developers and their customers. If what it takes to get standards set is some mysterious art practiced by people who wouldn’t recognize a byte if it bit them, then so be it. I just hope we get the DVD mess figured out soon. I have no intention of learning to troubleshoot DVD drives and media incompatibilities.

About the Author

Tim Mostad says, “the majority of my 18 years at Intel have been spent in the pursuit of technical marketing nirvana.” He is responsible for demos, white papers, plugfests and technical training to support the adoption of new desktop technologies.

Pentium® III Processor Platform Series

This new column provides the latest information on platform technologies arriving with the Intel® Pentium® III processor.

Streaming SIMD Extensions Revolutionize Human Interface Applications

By Shreekant Thakkar
Principal Processor Architect
Microprocessor Architecture and Planning Group
Intel Corporation

By Vladimir Pentkovski
Principal Engineer
Microprocessor Product Group
Intel Corporation

The Pentium® III processor introduces Streamed SIMD Extensions that enable software developers to enhance the performance human interface applications in volume PC market segments. The results will be visible improvements in 3D graphics and real-time video encoding, in addition to better and faster continuous speech recognition.

Principal Processor Architects
Today, with the introduction of the Pentium® III processor, Intel is launching Streamed SIMD Extensions to address the need for an increasingly important role for human interface applications in the volume PC markets. These extensions enable visually perceptible differences in 3D graphics, real-time video encoding, and higher quality, faster continuous speech recognition.

The new instruction architecture improves the performance of these and other Pentium III applications beyond what would have been possible from Moore's Law, and is scalable for future implementations.

Moving Beyond the Standard
Improving 3D graphics performance requires accelerating both front-end (geometry) and back-end (rendering) parts of the 3D graphics pipeline. Because floating-point (FP) computation is the heart of 3D geometry, speeding up FP computation is key to overall improvements in 3D performance. Rendering pipeline is done by a rendering accelerator, and a visually perceptible difference requires at least a 2X increase over the native FP performance in Pentium II processors.

3D graphics applications require the same computations as 3D data types (vertices), which makes them amenable to Single Instruction Multiple Data (SIMD) parallel computation models. This is the most cost-effective way of accelerating performance in 3D applications, and is similar to the acceleration for integer applications provided by MMX™ technology extensions. Thus an SIMD-FP model was selected for the Pentium III processor.

What does this Mean to Programmers?
This instruction set scope expanded traditional 3D geometry to include feedback from ISV usage of the MMX technology, as well as support for 3D software rendering, video encode and decode, and speech recognition. Streaming (prefetching and cacheability) instructions were added to increase concurrency between execution and memory accesses. In all, 70 new instructions and corresponding new state to IA-32 architecture were added. This is the first new state added since the 1386 processor added x87-FP

The programmer now has access to a rich set of SIMD-FP instructions, new media instructions, and streaming prefetch and cacheability instructions. These can be used simultaneously with MMX technology instructions to accelerate many media applications as well as non-media applications.

Benefits to Streamed SIMD Extensions
Programmers will see notable improvement in a variety of applications once the Streamed SIMD Extensions are utilized. While advancements in operating systems and databases will be noticeable, the performance increases will have the highest impact in these applications.

3D Geometry—SIMD-FP instructions can be used to improve performance of the 3D geometry pipeline to enable richer visual content and physical motion.

MPEG-2 Decode—the new instructions can be used simultaneously with MMX technology instructions to decode video, and with the new media instructions to decode audio.

Video Encoding—new media instructions can be used for soft VCR, video editing and enhancement of video conferencing applications.

Speech Recognition—the new media instructions can reduce error rate in recognition phase or increase vocabulary size.

Intel is constantly working to push the envelope and provide support for applications and performance. With the Streaming SIMD extensions, Intel® architecture accelerates performance for a different class of applications and breaks the barrier in expected growth for complex data sets.

About the Author
Shreekant (Ticky) Thakkar is a principal processor architect in the Microprocessor Architecture and Planning group (MAP), responsible for strategic planning. He lead the Internet Streaming SIMD extension development for Pentium III processor. Prior to that Ticky was responsible for the development of the Pentium® Pro Multi-processor.

Vladimir Pentkovski is a principal engineer in Microprocessor Product Group in Folsom. He was one of the architects in the core team whick defined Straming SIMD Extension of IA-32 architecture. Vladimir led the development of Pentium III processor architecture and performance analysis group. Previoiusly he led the development of compilers, software and programming leaguages for Elbrus multi-processor computers in Russia.

For More Information
For more detail on how Streamed SIMD Extensions improve performance in 3D and video, see the article “A Closer Look at SIMD Applications” in this issue of Platform Solutions News.

Reference manuals are available

Streaming SIMD Applications: A Closer Look

By James Abel
Software Development Engineering Manager
Intel Corporation

By Mike Bargeron
Software Performance Lab Development Engineer
Intel Corporation

By Mike Phlipot
Software Manager
Intel Corporation

The Pentium® III processor enables software developers to accelerate complex algorithms using new SIMD programming techniques. While programmers must learn the new instructions and identify how they can be applied to their unique software problems, the performance improvements can be significant enough to compensate for the initial learning curve.

Before the Pentium® III processor, Single Instruction/Multiple Data (SIMD) operations were limited to integer formats available in MMX™ technology. However, many graphics-oriented algorithms are based on floating point numbers and are restricted by the slower scalar implementation in Pentium II processors.

The graphics industry drive to design new applications, each running faster than its predecessors, created an immense programming challenge. Given these processing requirements, rich graphics and video effects found on high-end workstations and programs were not feasible. Until today.

With the introduction of the Streaming SIMD Extensions in the Pentium III processor, complex algorithms can now be accelerated using the new SIMD programming techniques. While programmers will be required to learn the new instructions and identify how they can be applied to their unique software problems, performance improvements can be significant enough to compensate for the initial learning curve.

How Streaming SIMD Extensions Work
Traditional integer operations utilized a single stream of data; Pentium III SIMD instructions take advantage of the symmetry in the algorithms to process four operations simultaneously. The “Streaming” part of the name refers to new instructions that quickly stream data into and out of the processor. For example, prefetching commands one part of the processor to load data, while another continues with computations. This processing overlap significantly minimizes waiting time between operations.

The Streaming SIMD Extensions consist of approximately 70 new SIMD floating points, SIMD integers and memory management instructions. In order to see how the Streaming SIMD Extensions benefit programming and data processing, let’s look at two of the applications which see the highest performance increase under this new approach.

3D Graphics
There are huge gains in utilizing the new SIMD instructions in 3D graphics. Performance gains can be doubled due to the ability to run the same operations on multiple pieces of data at one time.

While benefits from Streaming SIMD Extensions are significant in both Transform and Lighting, this example is limited to only the transform operation as shown in the following illustration.

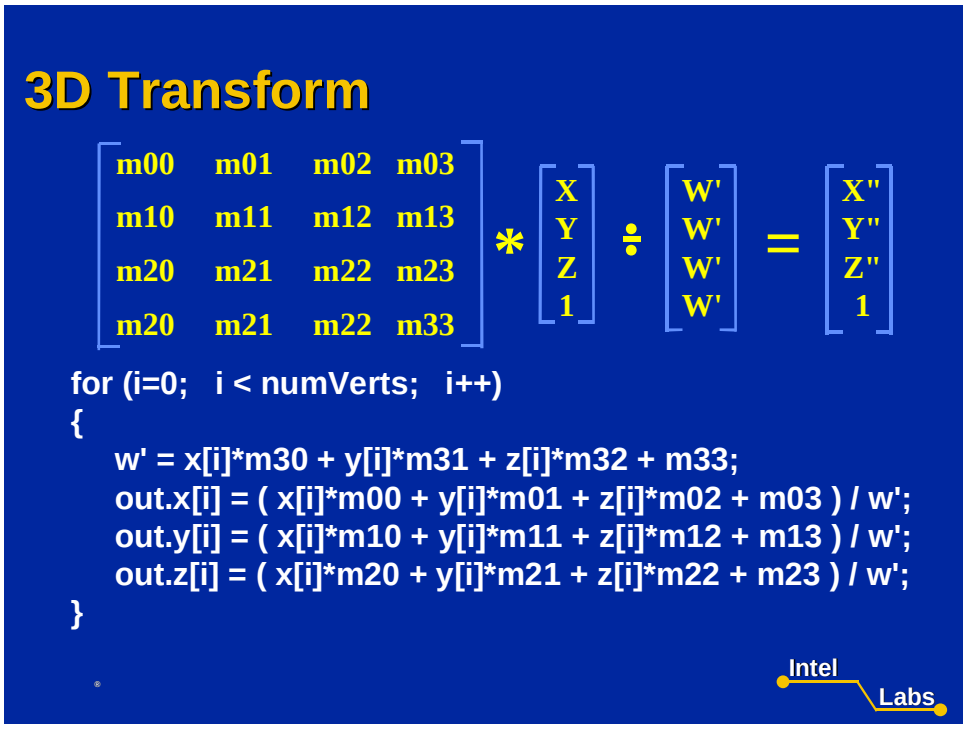

3D Transform

Recognizing that the same instructions will be performed on many different data elements allows the developer to take advantage of the Streaming SIMD Extensions in writing his/her 3D application. The following illustration shows how the Streaming SIMD Extensions can produce four times the number of results with the same number and type of instructions conventionally used in the 3D transform.

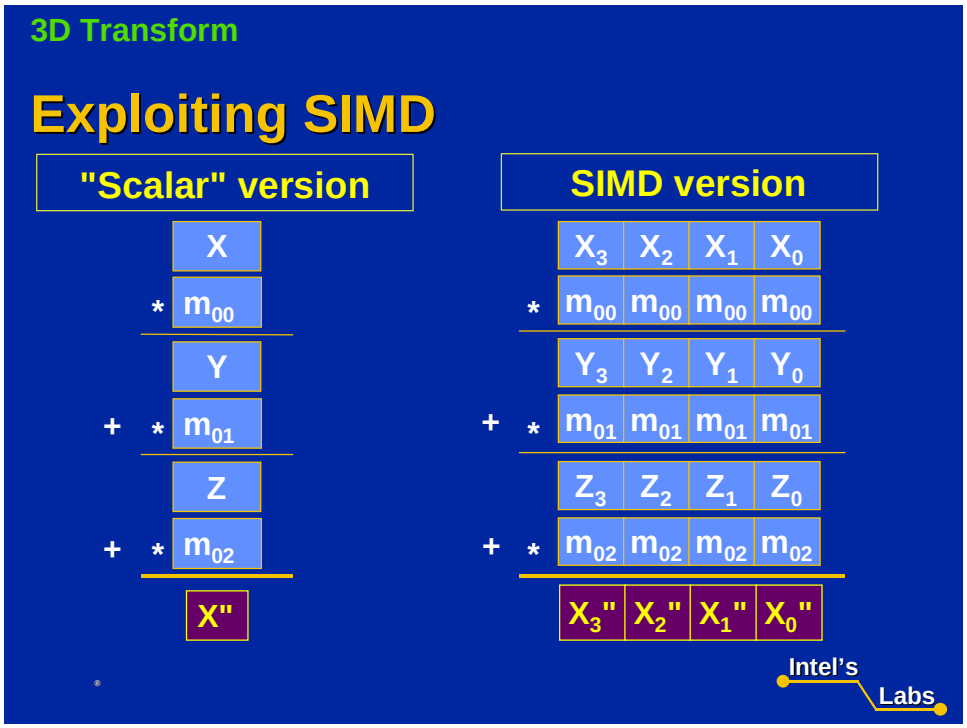

3D Transform – Exploiting SIMD

It is relatively simple to estimate which data elements will be needed for the next set of operations. “Prefetching” those elements allows the processor to request future data elements while it computes current results with data that is already available.

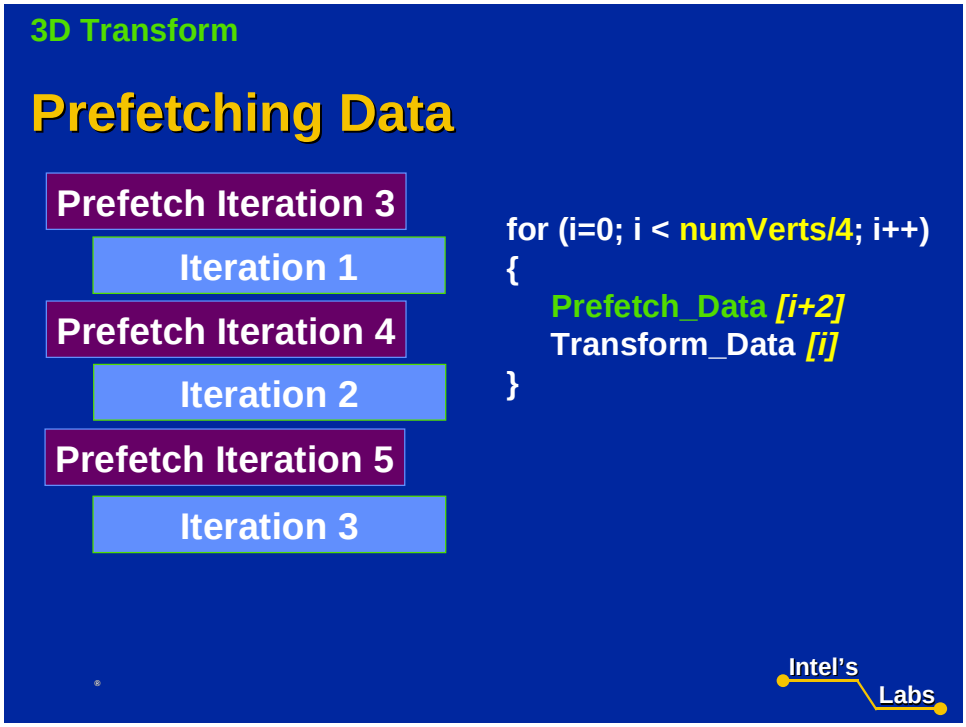

3D Transform Prefetching Data

Video Encoding and Decoding
Gains in video codecs are more incremental than in 3D graphics, but motion estimation is significantly faster than with just MMX technology integer instructions. The new Streaming SIMD Extensions include instructions which accelerate the subsampling and block matching operations—a critical part of motion estimation.

Block matching is the process of identifying the location (block) in the reference frame that is most similar to the block of interest in the current frame. A sum of absolute differences (SAD) computation is used to select the best match.

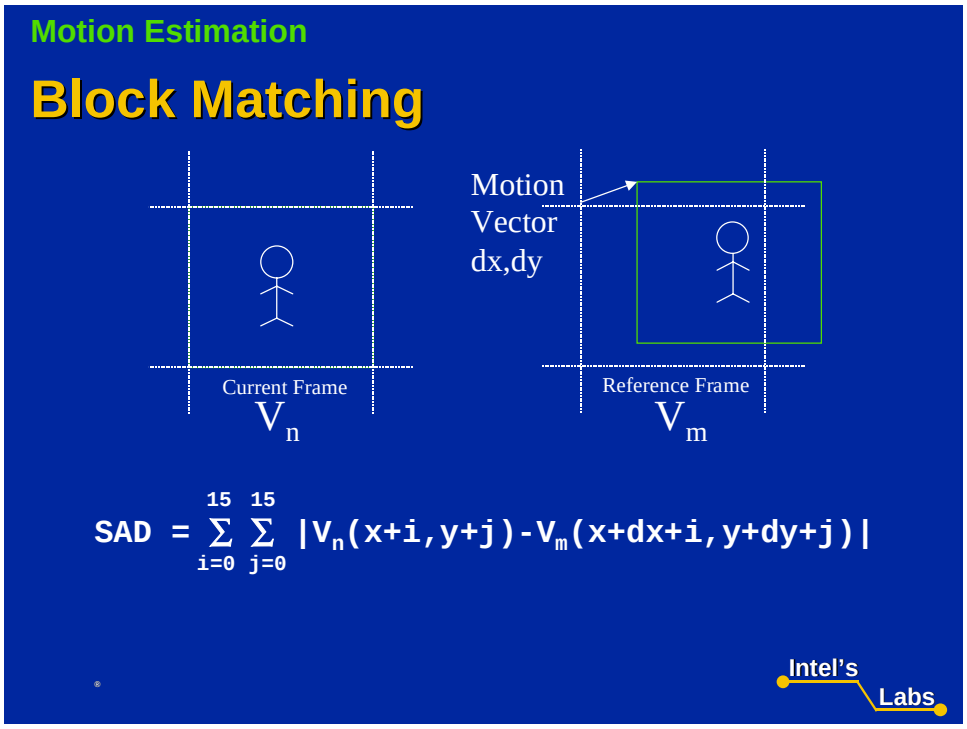

Motion Estimation – Block Matching

The new PSADBW instruction accelerates the block matching operation by performing eight absolute differences and a summation operation in one instruction.

The following illustration is an example of the block matching operation using the new PSADBW instruction. The result is an inner loop with significantly fewer instructions and a speedup of up to 200%.

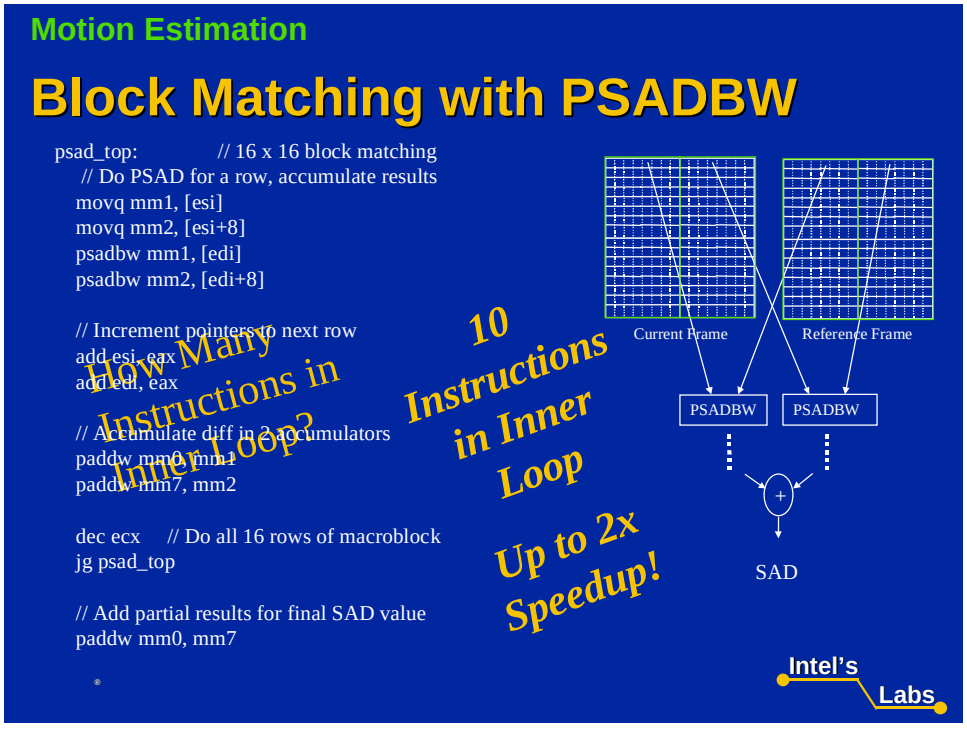

Motion Estimation – Block matching with PSADBW

The PAVG instructions can be used to speed up pel averaging. Use PAVGW for 16-bit words and PAVGB for bytes. This is used in a decoder (or local decoder in the case of an encoder) to perform interpolated motion compensation as may be encountered in B-pictures.

The Future of Programming Technology

With SIMD instructions based on integers with Pentium II and MMX technology applications, Intel led the drive to support graphics programs. The Streaming SIMD Extensions created for the Pentium III make it possible for richer graphics, higher quality, increased performance and faster processing. Programming instructions have moved into a new dimension….

About the Author

James Abel is currently a Software Development Engineering Manager in the Desktop Performance Lab at Intel Corporation in Chandler, Arizona. James obtained a Bachelor's Degree in Engineering from Bradley University in Peoria, Illinois in 1983 and a Master's Degree in Computer Science from Arizona State University in 1991. His interests include signal processing, computer architectures, software tools, and audio algorithms. James' email address is jabel@inside.intel.com.

Mike Bargeron started with Intel's Software Performance Lab in 1997. Since coming to Intel Mike has been involved in performance tuning 2D and 3D graphics applications for the PC. Specifically, he has worked with MPEG motion video as well as several 3D game titles. Prior to coming to Intel, Mike earned his BS in Electrical Engineering.

Mike works with desktop software developers to integrate Intel's newest processor capabilities into their applications. Most recently he has been helping 3D game developers take advantage of the Pentium III processor. In his ten years with Intel, he has held various engineering and management positions in technologies that include digital video compression and cable modems.

For More Information

For an overview of Streaming SIMD Extensions, see the accompanying article in this issue of *Platform Solutions News*.

Processor reference manuals and optimization manuals

Top Stories

The Merced Processor: Delivering Platform Solutions in 2000

By Hemant Dhulla
IA-64 Programs Manager
IA-64 Processor Division
Intel Corporation

Supported by partnerships with more than five dozen companies, Intel's Merced processor program is helping to ensure the availability of 64-bit Merced processor-based systems in the second half of next year.

Supported by partnerships with more than five dozen companies, Intel's Merced program team is working to ensure the availability of complete 64-bit Merced processor-based platform solutions in the second half of next year.

While Intel continues meeting its development milestones for the Merced processor—Intel's first 64-bit processor in the IA-64 product family, slated for production in mid-2000—equally important are its related development activities with a range of original equipment manufacturers and independent software and hardware vendors. That's because Intel's Merced program extends well beyond the development of the processor itself to ensure the availability of truly integral software and hardware platform solutions that will support and accelerate the introduction of Merced processor-based systems in the second half of next year.

As with all of its processor products, Intel recognizes the importance of working closely with platform developers to ensure that servers and workstations based on the Intel® Architecture (IA) provide users with the best possible performance and functionality. From enabling the creation of differentiated server and workstation designs to smoothing the development and tuning of tools, development environments, operating systems, drivers, and applications, the Merced program is paving the way for next year's introduction of systems that maximize the capabilities inherent in Intel's high-performance IA-64 architecture.

Mining Merced's Capabilities
While IA-32 processors will continue to provide ever-increasing performance for the mobile, desktop, workstation and server market segments, the next step forward for Intel and the industry is the Merced processor. Merced extends the capabilities of IA with scalability, high availability, 64-bit addressability, and the higher levels of performance and headroom required to meet the needs of the year 2000's most demanding workstation and server applications.

Merced provides 64-bit addressing, complemented by scalability features that include enhanced deferred transaction support to increase bus utilization, a large and innovative 3-level cache to reduce bus traffic, highly flexible page sizes to reduce virtual to physical address translation overhead, and speculation to reduce the effects of memory latency. For maximum availability, the Merced processor incorporates an enhanced Machine Check Architecture that coordinates error handling between the processor, firmware, and operating system, providing additional opportunities to correct and contain errors. And, the Merced processor includes innovative features such as data poisoning, which allows the corrupted data to be isolated and terminating only the affected process as opposed to the system, thereby increasing the overall system availability, and extensive parity and error correction circuitry (ECC). These features coupled with anticipated system features such as PCI hot plug and support from the industry's most highly available operating systems will allow Merced solutions to meet the availability requirements of the most demanding computing environments.

The performance of the Merced processor will provide solutions to meet the heaviest of computing needs. When the Merced processor is introduced, it will offer world class tpc and 3D graphics floating-point performance. Combined with its scalability, availability features and full IA-32 binary compatibility for end-user investment protection, this performance provides a uniquely compelling environment for the end users. Example deployment areas for end users in the server space include large memory databases, high-capacity OLTP, high-performance technical computing etc. kinds of applications and high-end software engineering, EDA, MCAD, DCC etc. kinds of workstation applications.

Merced at the Intel Developer Forum

Overall, there has been an unprecedented level of industry commitment to the Merced processor, with more than 60 leading original equipment manufacturers (OEMs), independent software vendors (ISVs), independent hardware vendors (IHVs), and operating systems vendors (OSVs) having announced their intention to provide Merced-based solutions. The February 1999 Intel Developer Forum (IDF) served as a key gathering place for Intel and its Merced platform development partners, validating the fact that:

- Merced processor development is on track for samples in 1999 and mid-2000 production,
- More than 30 OEM Merced designs ranging from 2 to 512 processor-based Server and Workstation systems are well underway and meeting their design milestones,
- Six industry-leading operating systems (Win64, SCO Unixware, HP-UX, Solaris, SGI IRIX and Novell Modesto) are booting on the Merced simulator,
- An extensive tools program is underway, including a compiler that is exceeding its performance targets, Software Development Kits (SDKs) that have been shipping to the industry to enable software development and an enhanced full development environment with a Merced simulator scheduled to ship to select ISVs later in Q1 '99, and
- A comprehensive server and workstation application optimization program is underway, creating the critical applications required for production environments with multiple server and workstation applications already running on Merced simulator.

Converging on Merced processor production in mid-2000, this industry support is poised to drive Merced solutions into production environments and establish IA-64 as a compelling architecture for the high-end server and workstation market segments.

Maintaining Momentum

The IA-64 track at IDF provided developers with a number of in-depth courses, including the Merced Solutions Overview, IA-64 Architecture Innovations, and Preparing 32-bit Code for IA-64. In concert with these activities, Intel announced that it will ship complete Software Development Kits (SDKs)—which will include a 64-bit operating system, compilers, libraries, and a comprehensive simulation environment—to select ISVs at the end of the first quarter.

For OEMs, ISVs, IHVs, and OSVs alike, the message is clear: the Merced processor is on track to make a significant impact in the computing world next year. As the first product implementation of Intel's IA-64 architecture, it provides the performance, scalability, and availability that today's high-end server and workstation markets require, extending the Intel architecture further into the high end of computing than ever before.

About the Author

Hemant Dhulla is a marketing manager in Intel's IA-64 Processor Division, where his responsibilities include developing and managing Intel's marketing and communications programs for the Merced 64-bit processor.

For More Information

For more information on Merced, IA-64, and associated technologies, please check out the following Web site: Merced Processor and IA-64 Architecture

The Web site is divided into three main sections. Extending the Intel Architecture provides details on the IA-64 roadmap and explains how IA-64 fits in with IA-32 to create a unified architecture. IA-64 Architecture Overview provides details on the innovations underlying the IA-64 architecture. And IA-64: The Unifying Architecture lists the more than 60 companies that have publicly committed to developing products for the Merced processor and IA-64 architecture.

Clarifying the Digital Display Interface Picture

By Steve Spina
Strategic Initiative Manager
Desktop Products Group, Intel Corporation

Intel and other core members of the Digital Display Working Group have created the Digital Visual Interface (DVI) specification to support the development of digital flat panel displays and help the industry move from VGA to pure digital technology.

Your first glance will convince you that digital displays provide superior resolution and greatly enhance the PC user experience. As the industry begins to move from analog to digital technologies, PC displays are poised for some dramatic changes. New space-saving and ergonomic flat-panel technologies are emerging, based on TFT LCDs and plasma technologies. At the same time, the display interface is beginning its transition from VGA to pure digital technology. Both of these trends promise to create significant new opportunities for developers.

The implementation of new flat-panel digital display products has been slowed by industry fragmentation caused by the multiple technical specifications that now exist. A universal and industry-wide digital display interface specification is essential to enable the more rapid and widespread adoption of these new display technologies.

While several technologies now exist that try to address the digital display interface issue, these are based on proprietary technologies and take different approaches. As a result, display interface technologies have not been open to all developers, and the goal of platform-independence has remained out of reach.

Join the Digital Display Working Group
At the September 1998 Intel Developer Forum, Intel led the formation of the Digital Display Working Group (DDWG), joining with other core members including Compaq, Fujitsu, Hewlett-Packard, IBM, NEC, and Silicon Image. The DDWG's mission is to remove the barriers to the implementation of digital display products by encouraging the involvement of developers from throughout the industry.

To meet this objective, the DDWG has created the Digital Visual Interface (DVI), a robust, comprehensive and extensible industry specification that defines the digital interface between digital displays and high-performance PCs. The DVI specification addresses protocol, electrical, and mechanical definitions and is the first specification expressly written for the Transition Minimized Differential Signaling (TMDS) digital interface, which ensures a degree of backward-compatibility with earlier approaches.

Implementation of products compliant with the DVI specification enables developers to immediately save costs by eliminating digital-to-analog conversion functions, while laying the groundwork for the eventual elimination of analog technology altogether. For PC users, digital displays provide sharper, more realistic visual experiences and automatically adjust to optimal resolution levels.

Royalty-Free Licensing
The DVI specification is based on "open intellectual property." This means that participating companies agree to royalty-free, reciprocal licensing of patents and other forms of intellectual property that are needed to implement the interface portion of the specification.

Intel and other DDWG members have made rapid progress on developing the DVI specification, culminating in the v0.9 final draft release of the specification at the February Intel Developer Forum.

Time to Get Involved
The DVI universal digital display interface specification enables the industry to deliver higher quality display products in innovative form factors at lower cost. These products will further enhance Visual Computing experiences and Ease-of-Use for users of Intel® Architecture PCs.
DDWG membership is open to all interested developers. Participating in the DDWG and designing DVI-compliant products will enable developers to deliver products that enhance the visual computing experience with PCs based on Pentium® II and Pentium® III processors.

The emergence of dramatic new digital display products is just beginning. Developers who are interested in becoming involved should take this opportunity to join Intel and other members of the DDWG. Participants can deliver input for the development of the DVI specification, and most importantly, they can use the specification to implement the design of DVI-compliant products.

About the Author
Steve Spina is Strategic Initiative Manager at Intel Corporation, specializing in graphics and digital display technologies that support Intel's Ease-of-Use and Visual Computing initiatives.

For More Information
For more information about the DVI specification and the benefits of DDWG membership, visit the DDWG Web site.

UDI and I$_2$O™: Complementary Approaches To Portable, High-Performance I/O

By Ramin Neshati
Technical Program Manager
Developer Guides for the Server Industry Marketing Group
Intel Corporation

Two complementary industry efforts, the Uniform Driver Interface (UDI) and the I$_2$O software specification for intelligent I/O, are designed to standardize I/O interfaces to support device driver portability, while improving performance.

Input and output devices on computers, such as storage or network adapters, are accessed and controlled by the operating system through device-specific software known as device drivers. In recent times, these I/O devices have become more intelligent and sophisticated and, consequently, their drivers have evolved from simple routines into highly complex, multi-layered programs. Because each operating system has its own unique I/O system and device interface, hardware vendors have been faced with the challenge of creating a different device driver for each operating system (including each variant of the UNIX* operating system) that they want their devices to support.

Two industry efforts are focused on standardizing platform and operating system I/O interfaces with the aim of increasing device driver portability, efficiency, and performance while reducing the maintenance burden for device driver developers. **Project UDI** is developing the Uniform Driver Interface primarily focused on the UNIX operating system vendors, and the **Intelligent (I$_2$O) Special Interest Group (SIG)** is creating a specification for intelligent I/O solutions. More information about UDI can be found on the Internet at the Uniform Driver Interface Web site. For further information about I$_2$O, visit the SIG Web site.

These two projects are complementary efforts that together will benefit the industry, business Information Technology (IT) departments and end users by promoting portable, high-performance drivers for the full range of I/O devices.

UDI: Portability Across Operating Systems
UDI is oriented toward a broad array of devices and has the goal of enabling the driver for any given device to work across different platforms and operating environments, particularly across various implementations of the UNIX operating system. This effort has yielded an architecture for device drivers that provides common interfaces to which they can be written. In effect, OS-to-driver communication has been abstracted to facilitate portability. While the driver source remains the same, recompilation for platform- and OS-specific considerations are sufficient to make that driver portable from one environment to the next.

UDI is developed through the cooperation of major industry forces such as SCO, Hewlett-Packard, Compaq (Digital), Adaptec, Interphase, IBM, Sun Microsystems, and other participating companies. Intel has recently joined this effort. The UDI specification is freely available to the industry and provides common interfaces for device driver developers on a wide range of system services, as well as interfaces for intra-driver communication. It facilitates driver encapsulation for portability across various hardware and operating system environments.

I$_2$O: Intelligent, Distributed I/O
I$_2$O technology, developed by the I$_2$O SIG, is a cornerstone of intelligent, high-performance distributed I/O for the Intel® architecture, including Intel's forthcoming 64-bit IA-64 processors. The objective of I$_2$O is to provide a uniform approach that complements existing drivers and offers a portable framework for the rapid development of a new generation of portable, I/O solutions.

I_2O is focused on intelligent, high-performance I/O subsystems, in which an I/O processor (IOP) offloads much of the work of controlling an I/O device. IOPs improve throughput by handling tasks such as buffering and transferring data and handling interrupts. I_2O supports message passing between multiple intelligent I/O processors.

Complementary Approaches
I_2O and UDI can easily coexist, giving hardware developers maximum portability for their devices and a chance to choose whichever interface best matches the characteristics of their device:

- UDI is the driver API specification of choice for non-intelligent devices (i.e., those that don't have an IOP), as well as for intelligent devices. It is also preferred for devices that stream their data.
- I_2O is the driver architecture of choice for intelligent storage and network devices with an IOP. It is also the logical choice for devices that would benefit from an IOP and for which the target operating system has an available OSM.

Models for Integrating UDI and I_2O
I_2O OSMs can be implemented as part of the UDI stack. Performance becomes a concern whenever a new layer is added to a software stack, but if the performance impact of running UDI as an extra layer is negligible, then running I_2O under UDI adds an extra dimension of portability. Along with allowing non-I_2O adapters and drivers to coexist with I_2O, it also ensures that older operating system module types can be upgraded across all operating system varieties, and that new module types can be developed and introduced by third parties.

UDI & I_2O: Win/Win for Developers and Users
The availability of both UDI and I_2O creates a win-win situation for hardware and operating system vendors as well as for businesses and end users.

Operating system vendors can reduce their development costs and ensure a wide range of device support for the operating systems without having to undertake the expense of developing the actual drivers.

Platform and hardware device vendors can create portable, easily maintainable drivers when they choose I_2O and UDI, while at the same time reducing their development and validation costs.

IT benefits by having improved interoperability and fewer and shorter qualification cycles for new devices and hardware. In addition, the availability of industry-wide driver and I/O specifications should work to increase the number, range, and variety of hardware devices.

Through its participation in the I_2O SIG and Project UDI, Intel will work with other industry leaders to create maximum synergy between I_2O and UDI and help realize the benefits of these two different yet complementary approaches to device portability and high-performance I/O.

For More Information
For more detail on I_2O and UDI, read the more comprehensive version of this article.

Visit the I_2O Web site for more information on I_2O.

Spurring Growth in the New Server Appliance Market Segment

By Lauri Minas
Manager
Server Industy Marketing
Intel Corporation

Intel is working with network and telecom industry leaders to develop a Network Server Appliance specification. The goal of the effort is to foster development of robust supporting applications, while creating opportunities for market growth.

Intel is working with network and telecom industry leaders to develop a Network Server Appliance specification that will define platform basic building blocks and interfaces. Learn how this specification will foster development of robust, supporting applications and create opportunities for volume market growth.

The rise of the Internet is creating extraordinary demand for new types of servers that can be deployed rapidly and managed remotely over the network. Voice carriers, telecommunication equipment providers, and Internet service providers (ISPs), as well as Information Technology (IT) departments at large businesses all need access to these new servers to keep pace with fast-moving technology and market changes, such as the convergence of voice and data. Reliability and simplicity are also key since server downtime can mean lost revenue and disrupted business processes.

Server appliances are fast emerging as the solution for highly reliable, low-maintenance servers. These fixed-function computers are designed to perform a specialized set of network-related server functions including print serving, web caching, web security, directory serving, and supporting search engines, groupware or data storage. Offering a complete hardware and software solution, server appliances run an operating system that is optimized for a specific function and are designed to be remotely managed over the network.

Overcoming Industry Fragmentation

The server appliance market segment is poised for explosive growth. Forrester Research predicts expansion from 91 million units in 1998 to 932 million in 2002. However, the market segment is beginning to fragment. Different device types and widely divergent approaches are evolving—all aimed at solving similar problems. This lack of a common approach could severely inhibit market segment growth. Without an industry common approach, software vendors hesitate to develop critical supporting applications and businesses tend to hold off purchases. In fact, Forrester Research's survey of Fortune 1,000 businesses found that 72 percent of respondents, while very interested in server appliances, are postponing purchases of these machines because of the lack of industry standards and specifications.

Industry Solution: Scope and Milestones

To address the issues of this new, fragmenting market segment, Intel and leading network and teleco vendors announced the formation of the Server Appliance Industry Group. Formed on December 7, 1998, this industry group is working on two key fronts: First, to define common architectural approaches and develop a specification that will incorporate existing industry defined standards. Second, to define a set of basic building blocks that simplify form factor as well as hardware and software interface designs for Intel-based server appliances. The specification will be for 'headless' and rack type server appliance platforms used by Telco vendors, ISP, and IT markets.

The specification created by the Server Appliance Industry Group will be OS-agnostic, and will provide the volume market opportunities necessary not only for ISVs to develop robust applications for Intel-based server appliances but also for appliance vendors to achieve economies of scale for their solutions. The specification will not be implementation specific, but rather will focus on capabilities. In addition to the specification, the industry group will also recommend sample usage models. Table 1 summarizes some key characteristics of the proposed Network Server Appliance specification.

Table 1. Network Server Appliances	
Sample Uses	File serving, caching, proxy serving, gateways, load balancing, firewalls, e-mail
Memory	Typically large, such as 64 MB with addressability of 4 GB or greater
Form factor	Thin form factor for rack-mount usage. “Headless” operation or no monitor attached
Operating System	Software architecture is optimized for the specific function of the device; not a full, multi-function operating system
Disk	Large capacity, such as 4 GB or greater
I/O	Fast network connections, such as 100-MB Ethernet and support for perhaps as many as 16 devices
Management	Remotely manageable over a network or Web connection
Other standards	Communications standards such as NEBS or ETSI

The Server Appliance Industry Group is sponsored by Intel, and the official charter and definition of this industry cooperative are expected in the first quarter of 1999. Key target milestones, including the release of the first specification, reference implementations, development tools, and test suites, are planned for the second quarter of 1999. Interested individuals and developers are encouraged to contact your Intel representative for meeting and specification development details.

About the Author

Lauri Minas is General Manager for Intel’s Server Industry Marketing. In her current position, she is responsible for the strategic direction of server industry efforts for Intel, and for marketing programs of server technologies across Intel divisions. Previously, she managed Intel's Wired for Management (WfM) Initiative, a cross-industry program to improve the manageability of business and home computers, which includes the WfM Baseline Specification. She has been at Intel for 16 years, holding various positions in marketing and management. Lauri Minas received an MBA degree in 1992, and a BA in 1979, both from Arizona State University.

Reducing Total Cost of Ownership with Intel® WfM and Microsoft ZAW Initiatives

By L. D. Weller
Manageability Marketing
Platform Marketing, Intel Desktop Products Group
Intel Corporation

When problems occur, PC users usually don't care whether it is hardware-related or software-related. That's why Intel and Microsoft are working together with complementary initiatives to improve manageability and lower total cost of ownership.

Bottom line, when computer users have a problem with their system, they want it solved immediately. They don't typically care if the problem is hardware- or software-related; they care about lost productivity, missed deadlines, getting their jobs done.

That's why it is advantageous for hardware vendors and software vendors to work together with complementary initiatives to drive PC hardware and software manageability technology.

Intel and Microsoft are doing just that by working together with complementary initiatives to lower the total cost of ownership for client systems. Intel's Wired for Management (WfM) and Microsoft's Zero Administration for Windows* (ZAW) initiatives provide this type of complementary technology, and include alignment on system-level information or instrumentation and pre-boot access methods.

WfM focuses on Intel® Architecture hardware and related system definition, and ZAW focuses on Windows* software environment. WfM and ZAW are separate, but complementary capabilities. One does not replace the other; rather, there are capabilities from each initiative that work together. For example, a WfM-enabled system that has an operating system taking advantage of the features of ZAW is more manageable than a system running either initiative separately.

For detailed information regarding Microsoft's Zero Administration for Windows (ZAW) initiative, please refer to Microsoft's Web sites and literature.

Overview of Wired for Management (WfM)
Intel's WfM initiative focuses on solutions to reduce computing ownership costs through better managed hardware platforms.

WfM is an Intel led initiative with broad industry support designed to provide common manageability capabilities for the purpose of reducing costs of ownership. Products based upon the initiative have been available and shipping since the first half of 1998.

The WfM initiative delivers the Baseline specification, and tools for both system and application vendors to adopt the technologies described in the Baseline spec. WfM recommends fundamental changes to the platform that are implemented by hardware system vendors.

For instance, in version 1.1 of WfM, the primary Baseline technologies are: Remote wake-up capabilities over the LAN; A universal network boot that allow LAN access to the device before it boots to its local resources (hard drive, floppy disk, etc.); Power management; and system level Instrumentation.

These technologies are either based on the work of standards bodies, such as the Desktop Management Task Force, or are incorporated in joint Intel/Microsoft design guides, such as PC99 or the Network PC Specification.

System vendors in 1999 are implementing these technologies into their devices and major software vendors are either shipping or have announced products that benefit from accessing WfM-based features.

A Complementary Approach

Although Intel and Microsoft approach manageability from unique viewpoints utilizing different products and initiatives, the end results are mutually beneficial solutions that provide new cost-saving capabilities to end users and IT departments.

Implementation of products at customer sites based upon these initiatives will prepare end-user organizations to be compliant with the manageability requirements outlined in the latest hardware design guides for Windows platforms (PC98, PC99, WHIIG, and the Server Design Guide).

Any organization looking to gain more control and prepare for future manageability capabilities should be investing resources today in the products offered based upon the manageability initiatives from Microsoft and Intel.

Intel and Microsoft are working together to add management value to the PC—to reduce the cost of ownership and to give IT greater ability to control the systems environment.

Following are key highlights of how the Intel and Microsoft initiatives work together to improve PC manageability:

- Within WfM there are elements, such as power management and universal network boot, that are synergistic with Microsoft's efforts on ACPI and the remote installation technology associated with Intellimirror*.
- Also, as instrumentation technologies have matured and grown within the Desktop Management Task Force (DMTF) standards' group, a new method of storing data—called the Common Information Model (CIM)—has emerged. CIM comprehends accessing information within the Desktop Management Interface (DMI) framework as well as within multiple other data storing methodologies.
- The Web-Based Enterprise Management (WBEM) initiative was turned over to the DMTF in April of 1998. BMC, Cisco, Compaq, Intel, and Microsoft announced the transition of this initiative to a standards-body in order to provide a vendor-neutral forum for a consolidation and unification point for manageability information stored in multiple existing manageability formats. These existing formats include SNMP, DMI, CMIP, and others.
- As a key part of WfM and ZAW, efforts are being made to provide a way to access DMI information and instrumentation from other sources through CIM.
- Microsoft's Windows Management Instrumentation (WMI) method of accessing data is based upon the object classes defined within CIM by the DMTF.

Recommendations

To best meet future computing needs and drive down total cost of ownership, hardware and software developers should:

- Specify WfM-enabled hardware platforms today.
- Specify use of ZAW technologies for managed Windows environments: Windows NT 4.0: Zero Admin Kit, Systems Management Server 1.2; Windows 2000 Professional and Windows 2000 Server: Intellimirror; and Systems Management Server 2.0.
- Specify and use WfM-enabled systems, and ZAW technologies together to reduce total cost of ownership and bring more control over the systems environment.

About the Author

L. D. Weller has been at Intel for 8 years and has been working with manageability products and initiatives for more than 10 years. He is also the chairman of the board of the Desktop Management Task Force, an independent, not-for-profit corporation providing standards and specifications.

For More Information

Visit the Wired for Management Web site.

For detailed information regarding Microsoft's Zero Administration for Windows (ZAW) initiative, please refer to Microsoft's Web sites and literature.

Home networks made easy—with no new wires

By David Redelfs
Technology Marketing Manager
Connected Home Initiative, Intel Architecture Labs
Intel Corporation

Through its Connected Home initiative, Intel is working with industry to help develop new applications and devices that deliver the computing power, connectivity and storage of the PC throughout the home, throughout the day.

Currently home networking is in its embryonic stage of development. Even homes with multiple PCs rarely have networks. The relatively few home networks that do exist are usually built by homeowners with technical expertise who typically create a miniature corporate network within their homes, complete with a hub, CAT 5 wiring, and ethernet adapters.

Today, home networking is still largely defined by LAN technology created for business users. Ethernet-based LAN configurations are too complex for most homeowners, and their installation requires pulling wire. As a result, today's home PCs are essentially islands unto themselves, locked in the den (or bedroom or other home venue) providing value only when users are within arm's length of the computer. To extend the value of the PC anywhere in the home, home networks are needed that are simple to install and easy for the homeowner to use. Saving the time, expense, disruption, and hassle of pulling wires requires "no new wires" networking.

Through its Connected Home initiative, the Intel Architecture Lab is working with industry to enhance the value of home computing by developing new applications and devices that deliver the computing power, connectivity and storage of the PC throughout the home, throughout the day. Imagine, for example, the ability to access personalized news, sports, and stock information from a wireless flat panel display while eating your morning oatmeal. Or, using a similar device in the family room to look up TV show listings, turn on and program the VCR, browse the Web or read e-mail while watching TV. Or how about being able to pick up a wireless phone, and through a voice command, have the PC look up the number via an Internet directory service and dial the call. These are just some of the benefits that can be realized when the power of the PC is freed from the den and made available where and when you want it via no new wires home networking.

There are five leading technologies for home networking requiring no new wires.

- **Home Phone Lines**
  In addition to carrying voice information, home phone lines can also be used to implement a data network for sharing peripherals, simultaneous Internet access through a single dial-up connection, file sharing, and multiplayer gaming. Intel is a founding member of the Home Phoneline Networking Alliance which developed the specification for home phoneline networking. Today, using a 7.5-MHz center frequency over telephone wiring, home phoneline networks can connect up to 25 PCs or other devices while supporting a data rate of 1 Mb/sec., while simultaneously supporting a normal phone call and an ADSL Internet connection. Planned future enhancements will extend the data rate to 10 Mb/sec. Intel has announced its involvement in the development of home phoneline networking products. See the January issue of Platform Solutions News for an update on Intel's recently announced 21145 Phoneline/Ethernet LAN controller.

- **Powerline**
  A typical home has an AC outlet in most every wall. And many home devices including PCs are already connected to the powerline. These factors make powerline a very convenient way to network devices in the home. For over a decade powerline has been used for home control networking, allowing a PC to control a home's lighting, heating/air conditioning, and security systems. In the past year, several companies have begun shipping, or have announced, powerline-based data networking products. These products claim data rates of up to 10 Mb/sec.

- **Coax**
  Cable TV wiring, while not as ubiquitous as either phonelines or powerline, provides another way for homeowners to network PCs and other devices. Cable modems are an increasingly popular means for PC's to attach at high speeds, to the Internet. And, coax provides an especially good way to deliver a PC's screen and audio content to an alternative output device, such as a TV, but can also be used for data networking as well. At least one company has a product that networks PCs, TVs, VCRs, gaming stations, and cameras over the home's TV cable, allowing PC applications to be run from any TV. This allows users to read e-mail from the kitchen TV, run a game on the large screen family room TV, or turn off the basement lights from the master bedroom TV.

- **HomeRF**
  Another approach to no new wires networking is to eliminate the wires altogether. Home RF (radio frequency) is a wireless infrastructure that provides portability and untethered freedom of movement anywhere in or around the home. Using 2.4-GHz frequency-hopping technology, RF signals can penetrate through walls and ceilings to provide a range of up to 150 feet. Intel is an active member of the HomeRF Working Group, which is working to establish an open industry specification for unlicensed RF digital communications for home PCs and consumer devices. Home RF supports both data networking as well as six channels of voice, with bandwidth dynamically allocated between the two.

- **Infrared**
  Infrared is another wire-free technology for home networking. Intel participates in the Infrared Data Association and was centrally involved in the development of the IrDA Control specification that defines a 75-Kb/sec. bidirectional link between a host device, such as a PC and a peripheral device such as keyboards, joysticks or mice. IrDA Control has a 30-foot range within a room.

  In addition to these no new wires transports, Intel is also involved in an effort that will make it easier for software programmers to monitor and control networked devices within the home.

Home API
The Home API initiative is an open industry effort to define a set of software services and APIs that enable applications to discover and control networked home devices such as TVs, VCRs, lights, security systems, thermostats, etc. It allows the software developer to issue a high-level command, such as "Set the TV to on," or "Set the VCR clock to 6:50 p.m." without worrying about either the underlying protocols or transport. By using Home API software, developers will be able to more rapidly develop home automation/control application.

Intel's Connected Home Initiative
Through its Connect.Home initiative, Intel is working to make its Anywhere in the Home vision a reality by extending the potential of home computing with new uses that deliver computing power and content throughout the home, throughout the day. This effort involves end-user research, the co-development of a variety of open industry specifications, the development of enabling technologies and new concepts for no new wires products and applications.

Developers are encouraged to participate in these emerging opportunities for no new wires home networking by becoming involved in one of the open industry specification development efforts, such as the HomeRF Working Group, the Home Phoneline networking alliance or the Home API Working Group. Intel is also interested in determining the level of industry interest in the formation of a home powerline networking and structured wiring initiative.

About the Author

David Redelfs is a Technology Marketing Manager for the Connect.Home Initiative at Intel Architecture Labs. He is the senior “marketect” and evangelist for IAL’s Anywhere in the Home vision.

For More Information

For information on the various home networking initiatives, please visit the following industry group Web sites:

- Home Phoneline Networking Alliance
- Home RF Working Group
- Infrared Data Association
- Home API Working Group

Read the white paper: Emerging Trends in Home Computing

Read Craig Hurst’s article Get Connected with Phoneline Home Networking in the January issue of Platform Solutions News.

Visit the Anywhere in the Home initiative page within the Intel Architecture Labs site. Intel Architecture labs site.

Alert on LAN 2: Remote Reboot to the Rescue*

Matthew Jung
Product Marketing Engineer
Network Communications Group
Intel Corporation

Alert on LAN 2 builds on the original Alert on LAN specification, providing IT managers with a way to reboot hung systems throughout the enterprise. It paves the way toward new levels of PC manageability.*

Intel has worked closely with IBM to deliver compelling new PC manageability technologies to the industry. The first example of this was Wake on LAN* technology—which today is a standard feature of managed PCs worldwide and comprehended by the PC '98, PC ' 99 and Wired for Management specifications. Shortly after the success of Wake on LAN technology, the Intel/IBM Advanced Manageability Alliance delivered Alert on LAN technology, which ventured beyond the Wired for Management (WfM) baseline.

The original Alert on LAN, developed in the Intel Architecture Labs, which was featured in the December 1998 issue of *Platform Solutions*, provides three principal manageability features—environmental alerts, a built-in timer and support for "presence heartbeat." Together, these capabilities for the first time have enabled managed PCs across the enterprise to send alerts to the corporate IT staff when problems such as a hung operating system or failure to boot have occurred.

But while the original Alert on LAN technology went a long way toward improving PC manageability by enabling IT administrators to quickly understand the nature and scope of problems, the thing it did not provide was a way for them to do anything to solve those problems. Enter Alert on LAN 2, the latest innovation from Intel and IBM, which allows IT organizations to remotely reboot systems after receiving OS hung or failure to boot related alerts—no matter where the system is physically located in the global enterprise network.

Return Path for Problem Resolution
Alert on LAN 2 technology's remote reboot feature provides today's platform developers with a compelling differentiating feature that can be employed in their high-end enterprise business PC product offerings. In addition, Alert on LAN 2 technology serves as a perfect complement to existing Alert on LAN implementations, which will continue to be available and see widespread application in the lower-performance segments of the corporate PC marketplace.

With Alert on LAN 2, IT administrators can remotely reboot a hung PC from central headquarters. And if the system still hangs, indicating a potential operating system problem, the staff can then boot the system up to a "known good state"—a pre-Windows** OS diagnostic state, for example. This enables IT organizations to circumvent a problem and run diagnostics.

Aside from serving as a powerful tool for desktop computers, Alert on LAN 2 is an ideal solution for mobile computers—many of which spend most of their time in a docking station, when not "on the road."

Building on a Strong Foundation
Alert on LAN 2 offers the same value-added features of its predecessor. Environmental alerts enable PCs to send alerts during the detection of chassis intrusion or the removal of the microprocessor, as well as alerts that correspond to such deviations as voltage irregularities and high chassis temperatures. As with Alert on LAN, environmental alerts can occur even if the PC is in a sleep state. A timer is built into the silicon which, if not refreshed—as in cases where the operating system is “hung”—automatically sends an alert out to the IT staff. Alert on LAN 2 also sends out periodic “heartbeats,” while the PC is in a sleep state, to indicate to the management console that the PC is indeed physically still connected to the network and that all is well.

Alert on LAN 2 for the Motherboard
OEMs can now begin to design Alert on LAN 2 with the Intel 82559 controller on the PC motherboard. The Alert on LAN 2 ASIC is slated for availability in the second quarter of 1999, while the i82559 is available today.

Alert on LAN 2 is the latest example of the combined efforts of Intel and IBM to advance the manageability of the PC platform. "IBM's Advanced Manageability Alliance with Intel has produced significant results, and Alert on LAN 2 demonstrates our commitment to reducing the cost and complexity of networked PCs through the development of innovative technologies," said Jan Janick, vice president, IBM Desktop Systems Development. "It all translates into better ways for IT managers to keep PCs up and running across the enterprise, and paves the way for customers to gain an advantage in today's highly competitive global marketplace.”

About the Author
Matthew Jung is a product marketing engineer in Intel’s Network Communications Group, where his responsibilities include working to define advanced LAN manageability solutions, Intel network controllers, and companion ASIC products.

*Wake-on-LAN and Alert-on-LAN are results of the Intel/IBM Advanced Manageability Alliance and trademarks of IBM.

**Third-party brands and names are the property of their respective owners.

For More Information
The Technology at Work for IT Managers web page has more about Alert on LAN 2 and associated technologies.

Intel® Mobile Processors: The Right Performance, Power and Packaging

By Kapil Wadhera
Business Development Manager
Mobile/Handheld Products Group
Intel Corporation

Intel® Pentium® II and Celeron™ processors support the higher performance, smaller form factors and longer battery life required by the next generation of mobile PCs.

As society becomes increasingly mobile, there are few products on the market that hold as much appeal, or as much promise, as notebook PC systems. Significant growth in sales over the last few years and recent trends are just the tip of the iceberg.

IDC and Dataquest, two independent research firms, are predicting healthy growth for the notebook PC market well into the next millennium. Aside from portability, there are a number of reasons for the increased popularity.

Notebooks are more powerful. The notebook processor performance increase has continued to accelerate–almost doubling in the last year alone. Today's mobile Pentium® II processors offer desktop performance in a mobile PC enabling new users and demanding applications.

Notebooks are more affordable. The notebook system prices are declining rapidly. Less than 18 months ago, notebook systems with the newest processor technology were sold for more than $5,000. Today, notebook systems with the latest technology are available for less than $3,000. In addition, the cost of owning a mobile PC is declining and fast approaching that of desktops.

Notebooks offer greater flexibility. Employees who need to work remotely can do so more easily. The improved access notebooks offer these workers can increase productivity and even employee satisfaction. This is one of the reasons why more and more businesses are converting from desktop PCs to notebooks. For example, Peoplesoft has converted 100 percent of its workforce to notebooks.

The selection is better than ever. Notebooks are available in multiple form factors including Full Sized, Thin and Light, and Mini-Notebooks systems. This wide variety of choices in mobile systems support the broad spectrum of users and needs.

The Technological Challenges
With the notebook industry on the verge of dramatic growth, Intel continues to focus considerable resources on developing new technology to support the needs of an industry driven to produce high performance and thinner notebooks.

While desktop computers can run high-demand applications using in excess of 30 watts of power, notebook systems require processors with significantly lower thermal dissipation and more sophisticated system technologies to maintain desired battery life. The challenge is to enable high-performance, thermally efficient processors in mobile friendly form factors.

The Solution: New Generation Pentium II and Celeron™ Processors
Intel recently introduced a family of new processors specifically designed to meet the high demands of mobile PCs. The new Pentium II and Celeron processors offer high performance for productivity, multimedia, and Internet applications. These processors are based on Intel's latest P6 micro-architecture and provide support for manageability and special low power modes.

The new mobile processors are the first Pentium II processors offered on a single silicon die. By integrating processor and the cache on a single die, mobile users benefit from higher performance and lower power consumption.

In addition, Intel has developed an innovative mobile processor packaging technology with a smaller, thinner, and lighter Ball Grid Array (BGA) package. The new BGA package is 86 percent smaller and 89 percent lighter than its mini-cartridge predecessor launched a year ago.

Pentium II for High-Performance Applications
The new mobile Pentium II processor is available in speeds up to 366 MHz and uses 256K of full processor speed, on-die second level cache.

The mobile Pentium II processor at 366 MHz delivers up to 27 percent higher performance than existing Pentium II processor at 300 MHz with off-die cache. In addition, the single-die integration provides higher performance without sacrificing battery life. In fact, power consumption on the new Pentium II processor at 300 MHz is approximately 15 percent less than the previous Pentium II processor at 300 MHz.

Mobile Pentium II processor based systems offer desktop performance in a mobile PC and are ready for today's demanding software applications and tomorrow's advanced operating systems.

Mobile Celeron Processors Offer Great Performance at an Exceptional Value
The new mobile Celeron processor offers all the benefits of increased performance and lower power at an affordable price. The Celeron processor provides 128-KB on-die level 2 cache, and is capable of handling everything including business, educational, and multimedia applications.

The Intel® mobile Celeron processors are available at 266 and 300 MHz. The Intel Celeron processor at 300 MHz delivers up to 35 percent performance improvement than the Pentium® processor with MMX™ technology at 300 MHz.

The Definitive Answer
Intel is committed to developing power friendly, higher performance, smaller form factor processors to meet the demands of the mobile market segment. All of this without compromising the quality and reliability end users and OEMs expect from Intel.

If higher integration, lower power and higher performance aren't enough, here's more: OEMs can offer all this with minimal re-engineering. New mobile processors are offered in all the same packages as the prior generation and utilize the same power supply voltages.

Beginning January 25, 1999, IBM, Compaq, NEC, Fujitsu, Dell, Gateway, Acer, and other leading mobile PC manufacturers launched their new line of mobile PCs with the single die processors. The new Intel® mobile processors deliver–higher performance, lower power, and smaller packaging–in addition to higher megahertz in mobile PCs. It's the solution that works best for all segments of the mobile PC market.

About the Author
Kapil Wadhera earned his Masters in Electrical Engineering from Arizona State University, and his Masters in Business Administration from the University of Santa Clara. With more than nine years' experience at Intel, Kapil has worked in MHPG marketing since 1994. Kapil was most recently responsible for managing the mobile Pentium II processor product line.

For More Information
More information about the new mobile Pentium II and Celeron processors, is available from Intel's Web site.

Delivering High-Value Content Through Digital Transmission Content Protection

By Michael Moradzadeh
Director, Strategic Planning and Architecture
Home Products Group
Intel Corporation

Many misconceptions exist regarding 5C Digital Transmission Licensing Administrator (DTLA) policies, but the truth is out there. Here are the facts.

Intel continually invests in delivering the value of convergence to our customers. One focus of our work is bringing high-interest Hollywood movie and music content to the digital world of the PC. In the past, we've worked with major studios, CE manufacturers, and our customers to assure that DVD movies are playable on PC's using software or hardware players.

Now, we're developing and delivering Digital Transmission Content Protection (DTCP) that lets a PC or digital TV connect to streaming, digital movies and other content from digital appliances, like set-top boxes, DVCRs, satellite receivers, and DVD players.

The 5C model of DTCP was developed by Intel, Hitachi*, Matsushita*, Sony*, and Toshiba*, hence the name, 5C (for "5 Companies"). DTCP uses an encryption scheme together with signed certificates to provide protection from copyright infringement—all while being transparent to the user. A wealth of information about DTCP can be found at its Web site. Visit the DTLA or Digital Transmission Licensing Administrator Web site, where you'll find the license policies, FAQs, and more.

This technology has been adopted by CableLabs* for the OpenCable specification. Products using DTCP to deliver movies between devices have now been demonstrated in Hollywood and was showcased at the Intel Developer Forum on February 23-25. A number of vendors are developing implementations for Intel-based systems, and prototypes of these have been demonstrated in Hollywood as well.

Clarifying Some Points
As DTCP is more broadly accepted, some of our customers (and even our competitors) have suggested areas for improving the way we manage the licensing of this technology. In response to these suggestions, the 5C has upgraded the license terms and has issued a policy memo expressing those changes.

Onward and Upward
At Intel, we know that our position in the PC and CE industries is due in large measure to our relationships with other corporate leaders. We work to enhance these relationships with other industry resources and the developer community in every way we can—every step of the way.

About the Author
Michael Moradzadeh is Director of Strategic Planning and Architecture in Intel's Home Products Group. He is very active in leading Intel's efforts in content protection and delivery.

For More Information
Digital Transmission Licensing Administrator Web site

Maxtor Measures Success with Intel® Storage Performance Toolkit

By Brian Johnston
IPEAK Marketing Manager
Platform Tools Organization
Intel Corporation

By JJ Clark
Performance Analyst/Technical Marketing
Maxtor Corporation

Despite their phenomenal gains in capacity over the last decade, disk drive speed improvements have been harder to come by. Intel's IPEAK Storage Performance Toolkit helped Maxtor designers boost product performance.

Disk Drive Manufacturer Reaps Big Rewards with Comprehensive Image-Based Performance Analysis Tools

When you're building a product that's highly subject to the physical limitations of materials and moving parts, performance gains come only with a great deal of effort. Take disk drives, for example. Despite their phenomenal gains in capacity over the past decade, speed improvements have been harder to come by. Today, disk drive manufacturers realize that drive speed is playing a growing role in the way users perceive overall PC performance.

The people at Maxtor Corporation know this perhaps better than anyone else, because they've been making disk drives since 1982. Not long ago they launched an ambitious effort to boost product performance, and with the help of the IPEAK Storage Performance Toolkit from Intel Architecture Labs Scalable Platforms Initiative, they did just that. Relying heavily on the toolkit for performance analysis, they introduced a brand-new architectural platform as well as their DiamondMax 4320 4.3 GB/platter, which Maxtor estimates is significantly faster than its competition.

They also chalked up impressive numbers on the balance sheet, according to Staff Engineer Jim Deng. "With IPEAK, the cost of licensing, implementing, training, and maintenance is insignificant compared to our savings in development time and costs, our boost in sales and our growth in customer satisfaction" he reports. "For a relatively small investment, we're reaping very large financial returns."

Traditional Benchmarks Limited

Principal Engineer Bob Small, explains, "With many benchmarks we would get a number indicating a certain performance level, but we didn't know what was going on underneath. We could learn *whether* we were faster or slower than a given competitor, but if we wanted to learn *why*, we had to hand-script numerous detailed tests using logic analyzers and other complex equipment."

The Market Doesn't Wait

As Executive Director of Firmware Development David Hunter explains, relying exclusively on scripted testing and logic analyzers for low-level performance evaluation is not only tedious but also time-consuming—in an industry notoriously unforgiving of delays.

With this in mind, Maxtor executives began talking with Intel about the IPEAK tool set. In response, Intel consultants worked closely with Maxtor engineers to show them precisely what the tool set could do for their disk drives. With one test, for example, the engineers discovered that many requests were requiring two revolutions to be serviced and taking two or three times as long as they should have to complete. "That was a big red flag for us," says Hunter. "We had no idea this was a problem when we were just running traditional benchmarks."

High-Level Focus, Low-Level Detail
Since then, IPEAK has played a central role in development at Maxtor, transforming the way firmware and hardware engineers do their jobs. “IPEAK keeps us focused at a high level even while providing us comprehensive coverage of low-level information,” says Hunter. “Now, instead of having to sift through thousands of lines of code looking for the source of a problem, we can devote our engineering resources to fixing it.”

For example, one new product Small was testing showed a significantly lower performance benchmark than that of a certain competitor. Initially he suspected the problem lay in the basic caching scheme. But an IPEAK test measuring read service time distribution showed him that the caching was as good if not better than the competitor’s.

This led Small to the actual source of the problem: the basic command overhead. “With IPEAK I quickly isolated the problem and began working toward a solution instead of chasing a dead end,” he explains. “This saved me weeks of work and helped us bring the product to market that much faster.”

In another case, Small wanted to confirm whether Maxtor was maintaining its lead in write-caching performance. Using an IPEAK test measuring the distribution of service times vs. seek distance for back-to-back writes, he learned that it did.

For Maxtor, this means that analysis of a caching scheme that used to take several days now can be done in a single afternoon. Other analysis tasks are equally streamlined. As Deng says, “With IPEAK, a firmware engineer can do most testing in about a tenth of the time it would take using a logic analyzer.”

Also Useful for Adding New Features
IPEAK also is helping Maxtor implement new features. For example, using the IPEAK test for distribution of service time vs. seek distance, engineers determined that they could achieve higher performance by implementing a form of zero latency read. (This is a process in which the drive head begins reading data from a location preceding that of the data requested by the host.) They did this by designing a new caching scheme and then confirmed its effectiveness by using the IPEAK test for overall distribution of service times. The result was a 15 percent gain in WinBench* 98 read performance.

See What Winning Looks Like
Whether Maxtor engineers are doing analysis or tracing, whether they’re identifying problems, implementing new features, or evaluating performance vis-a-vis the competition, IPEAK provides them the images that show exactly what’s going on inside their disk drives.

About the Author
Brian Johnston has over 23 years experience in the computer industry. He currently manages the team responsible for marketing Intel's IPEAK family of platform integration tools. In seven years at Intel, Brian has introduced software quality standards and helped enable industry adoption of new platform technologies.

JJ Clark is a Performance Analyst in the Technical Marketing Organization within the Maxtor Corporation. She coordinates competitive analysis, performance modeling and works along side team members in product launches. JJ has been in the storage industry for over 8 years and has extensive background in customer support as well as failure analysis.

For More Information
Visit Maxtor Web site.

See the IPEAK Web site.

Design for Compliance: PCI Test Metrics for PC 99

By Murthi Nanja
Staff Software Engineer
Intel Platform Compliance Operation
Intel Corporation

The Design for Compliance series continues this month with a look at the Peripheral Component Interconnect (PCI) chapter of the PC 99 Hardware Test Specification, providing objective pass/fail metrics for hardware elements of the PC 99 System Design Guide.

In the January *Design for Compliance: New PC 99 Test Metrics* article, we described how Intel Corporation and Microsoft Corporation are extending their PC 99 System Design Guide co-development effort and gathering industry input to create a new set of PC 99 Test Specifications and tests.

The *PC 99 Hardware Test Specification* provides a clear set of objective pass/fail compliance metrics for hardware elements of the *PC 99 System Design Guide*. This month's article showcases the chapter entitled, *Peripheral Component Interconnect* (PCI) from *the PC 99 Hardware Test Specification*.

Tests and PCI Test Assertions
Test assertions are declarations of truth statements derived from guidelines, industry standards and specifications referenced in the *PC 99 System Design Guide*. Test assertions are clearly stated in the *PC 99 Hardware Test Specification*, with reference to the source documents, and mapped directly to the tests. The PCI test specification contains nearly 200 assertions, and of those, over half can be tested through software. The complete list of assertions is organized into seven tests, each of which spans multiple test assertions.

The PCI compliance tests verify the hardware assertions defined in the *PC 99 Hardware Test Specification*. As designed, this feature makes the PCI tests independent of the operating system itself.

Sample PCI Test Assertions
Chapter 9 of the *PC 99 System Design Guide* identifies the requirements for PCI host controllers and peripherals. The chapter specifies that the system-board bus and all components on the system bus, must implement the *PCI specification*, Revision 2.1 or later. It further specifies that the PCI bus, any PCI-to-PCI bridges, and all add-on capable devices must comply with the *PCI Bus Power Management Interface Specification,* Revision 1.1 or later.

Requirements in the *PCI Power Management Specification* plus the ECNs are the source of the PCI assertions in the *PC 99 Hardware Test Specification*. A sampling of the assertions generated from the requirements include:

- PME# enabled by setting the PME_En bit in the PMCSR register.
- Once the OME# signal is asserted, the device must continue to drive the signal low until software explicitly clears the PME_Status bit in the PMCSR register.
- PME# must generate some form of SCI and system software must be responsible for handling the PME# service routine.

Clear Results
The test case output log begins with an initialization section and concludes with a summary. The initialization section identifies the test parameters, test start date and time, and other relevant target platform information. A summary section identifies test case pass/fail status.

To support failure analysis and troubleshooting, the tests have a verbose debug output mode. Depending on the debug level specified, the tests generate increasing levels of detail in the test results log files.

PCI Test Hardware Requirements
The majority of PCI hardware compliance tests require a special PCI-PM compliant test adapter. The test adapters are based on the PCISIG card and provide additional functionality supporting PCI power management requirements. Some tests require several adapters. The test adapters have the following power management features:

1. PME# assertion support from D0, D1, D2, D3 Hot, and D3 Cold
2. PME# signal assertion via software control
3. Programmable 3.3 Vaux load select (20 mA, 160 mA, or 375 mA)
4. LED indicators for various PCI voltages

In addition to the PCI test adapter requirement, most PCI tests also require ACPI-compliant target systems with S3 power-state. More information about the PCI test adapter can be found in the *SPECS* area of the System Test-IF Web site.

We Need Your Help!
The PCI test assertions contained in the *PC 99 Hardware Test Specification* are designed to provide objective PCI tests criteria that make it easier for developers to design PCI system-board and peripheral products. You now have an opportunity to download the latest test specification draft and provide your feedback on the PCI test assertion coverage and test algorithms.

The current *PC 99 Hardware Test Specification* plan targets a release of draft 0.7 in early March 1999, with the goal of releasing version 1.0 in April 1999. The Alpha version of the hardware tests will be available to System Test-IF members in early March 1999, while the public Beta release is targeted for early April.

We encourage you to review the latest Test Specification draft and evaluate the early versions of the compliance tests using your own products. If you discover any issues or just want to send us your comments, please visit the System Test-IF Web site and let us know. We're listening!

For More Information
Download the latest release of the *PC 99 Hardware Test Specification* and learn more about the System Test Implementers' Forum from the System Test-IF Web site.

About the Author
Murthi Nanja is a staff Software Engineer with Intel's Platform Compliance Operation, responsible for developing hardware validation tools for PCI aspects of the PC 99 System Design Guide. He has over 10 years of experience at Intel working on hardware-assisted debuggers, and system validation of Intel® processors and chipsets.

Technology News Bytes

Intel Launches the Pentium® III Processor

Intel has introduced the Pentium® III processor, the first microprocessor designed to power a new Internet experience filled with rich audio, video, animations and 3-D that make information come alive. The Pentium III processor is available immediately at 450 and 500 MHz, while the 550 MHz version will be available in second quarter 1999. The Pentium III processor at 500 MHz is 93 percent faster than the Pentium II processor at 450 MHz on CPU intensive 3-D calculations, as shown by Ziff-Davis' 3D WinBench* 99 transform and lighting test. Using Future Reality's MultimediaMark* multimedia performance benchmark, the Pentium III processor at 500 MHz is 42 percent faster than the 450 MHz Pentium II processor. For more information on Pentium III processor performance visit http://www.intel.com/procs/perf.

Intel Introduces VTune™ 4.0 Performance Enhancement Environment with Support for Pentium® III Processor

With the VTune™ CD, Intel offers a suite of performance tools to help developers extract maximum performance from their code on Intel® Architecture. The suite consists of: VTune tool, the Intel C/C++ Compiler and Intel FORTRAN Compiler plug-ins, and the Intel Performance Library Suite. It also contains the Intel Architecture Tutorials.

Smart Toys Learn New Tricks

On February 3, 1999, Mattel, Inc. and Intel Corporation announced the first two products in the Intel® Play™ line of PC-enhanced toys. These products were shown at the American International Toy Fair in New York on Feb 8-12. The products combine the power of PCs and the "magic" of toys into a new generation of interactive play. The Intel Play X3™* Microscope and Intel Play Me2Cam™*, will be available from Mattel in the fall of 1999, at an approximate retail price of $99 each. The products have been jointly designed and developed by a team of engineers and toy designers from the two companies.

Software Development Tools Announced for the Intel® Pentium® III Processor

Power for the Next Generation of the Internet and Desktop: Assemble a complete development environment for the Pentium® III processor. Extend the power of the Pentium III processor to Web Development and Internet technologies. Enable your software, content, and web sites to take advantage of the newest Intel processor by using the tools featured on this site.

Intel® Indeo® Video Improves with Version 5.1

The latest version of Indeo video, version 5.1, has been released by Intel and is now available for download. Version 5.1 provides a number of new installation improvements such as a convenient all-in-one installation package and an option for minimal video file playback installation. In addition, InstallShield* scripts are now provided to allow developers to customize their installation program to setup Indeo video software to meet their needs.

Intel® 3D RSX Licensed to RAD

Intel announces a license agreement between Intel and RAD Games Tools for the Intel 3D Realistic Sound Experience (3D RSX) software. RAD plans to incorporate 3D RSX into their Miles Sound System* product and extend the features and functions to continue to meet evolving customer needs.

Initially developed within the IAL Internet Media Initiative, 3D RSX offers a Dolby*-certified surround sound software library, which allows you to create rich sound experiences your users can enjoy without expensive add-in cards or custom speakers.

G.Lite to Enrich Asia Pacific Homes

Asia-Pacific consumers may soon enjoy less expensive high-performance Internet access. On January 21, 1999, industry leaders SingTel Magix and Intel Corporation, in conjunction with Newbridge Networks and Efficient Networks, announced they are conducting a field trial in Singapore to examine the feasibility of introducing a new mass-market version of Asymmetric Digital Subscriber Line (ADSL). Preliminary evidence indicates that this new ADSL technology performs according to specifications outside North America. The new ADSL technology, G.Lite, previously known as Universal ADSL, has the capacity to provide a cost-effective and standards-based solution for high-speed network access to homes.

EE Times Spotlights Intelligent I/O

In a series of articles, EE Times reviews the state-of-the art of intelligent I/O. First, look under the hood to see how the advanced data-flow architecture of Intel's latest i960® I/O processors improves throughput and server performance. Then check out the growing appeal of intelligent I/O solutions for mid-range and entry-level servers. Finally, discover how to choose the right design path for your application-raw I/O processor speed-or "hardware-optimized" I/O features.

i960® RN I/O Processor Powers Cyclone Compact PCI Controller

Optimizing I/O bandwidth is a key issue at Cyclone, and Intel's new i960® RN I/O processor provides the solution. Cyclone has implemented the new i960 RN I/O processor in its CPCI-943 Compact PCI Intelligent I/O Controller. Cyclone's design goal for the CPCI-943 was to implement high bus bandwidth, to support I2O™ (Intelligent I/O) technology, and to integrate 64-bit bus interfaces capable of unlocking potential bus bandwidth bottlenecks in embedded real-time systems. The CPCI-943 controller features dual 64-bit PCI buses and 528 Mbytes/sec SDRAM bus bandwidth and supports two PCI mezzanine (PMC) cards. The PMC-based design was chosen to allow customers to flexibly meet a variety of embedded I/O requirements without custom engineering, by adding various combinations of modules to the card.

Analog Devices and Intel to Develop Digital Signal Processing Technology

On February 3, 1999, Intel Corporation and Analog Devices, Inc. announced a joint development agreement to design a digital signal processor (DSP) core architecture. The companies will develop a fixed-point, low-power DSP core ideal for processing video, image, voice, and data in emerging embedded communication and computing devices.

DRAM RIMM Module Reference Designs

The memory bandwidth requirements of the mainstream performance PC will grow beyond the capabilities of a PC100 SDRAM in 1999. Intel has been working with Rambus and the DRAM industry on a faster main memory implementation for 1999 platforms. This new technology, known as Direct RDRAM, will provide peak data transfers rates of 1.6GB/sec. Direct RDRAM components come in uBGA packages and are assembled on RIMM modules.

Industry Events

Game Developers' Conference (GDC)
March 15–19, San Jose, CA
This is the annual industry event for developers who make entertainment software across all computing platforms.

WinHEC 99 (Windows* Hardware Engineering Conference)
April 7–9, Los Angeles, CA
Steve Ballmer, David Cole, Brian Valentine, and Carl Stork will be just a few of the speakers at WinHEC 99, Microsoft's annual gathering for the computer hardware industry. This year, Microsoft expects 4,000 managers, developers, planners, and engineers from semiconductor makers, peripheral makers, and system manufacturers.

—End of Platform Solutions Issue 17—